Be Prepared for PHY and PCIe® Controller Integration

Sampath Banka

Senior Applications Engineer

PLDA

Disclaimer

<u>Presentation Disclaimer</u>: All opinions, judgments, recommendations, etc. that are presented herein are the opinions of the presenter of the material and do not necessarily reflect the opinions of the PCI-SIG®.

Summary

- **Introduction**
- **Scenarios and Source of Issues**
- **Issues and Solutions During PCI Express® PHY - Controller Integration**
- **Conclusion**

Introduction

- Designers are skeptical while doing PHY & MAC integration, often consider it risky and prefer to avoid it, however if you consider the common scenarios for source of issues, you reduce the risk to a great extent.

How do MAC – PHY Communicate

- Preview of MAC and PHY Communication

  - Rx Detect
  - Rate Change
  - Equalization
  - Power Management
  - Data Sequence

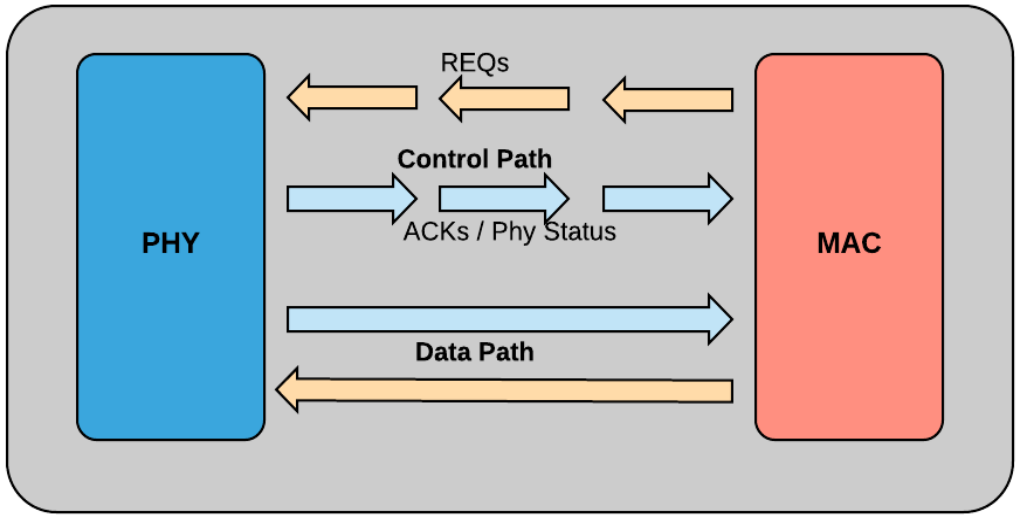

- Introduction
- **Scenarios and Source of Issues**
- Issues and Solutions During PCI Express PHY - Controller Integration
- Conclusion

Scenarios and Source of Issues

- During PHY & MAC integration, you would encounter the following scenarios:
  - Proprietary Behavior:
    - Lane-Lane Deskew
    - RxValid Filter
    - N_FTS
  - PHY Clock generation
  - Low power L1PM integration
  - Lane Margining Implementation
  - Equalization Issues
  - PHY Specific Options
  - Architecture Challenges

- Introduction
- Scenarios and Source of Issues
- **Issues and Solutions During PCI Express PHY - Controller Integration**
- Conclusion

Lane-Lane De-skew

- PCIe® Base Spec provides the Guidelines for the MAC to support the Max possible Skew in real hardware.
- But some PHY SERDES worst case Lane-Lane skew exceeds these numbers.
- When MAC doesn't support the range, this causes symbols to lose alignment across lanes.
- This effect is more visible at 8.0 & 16.0 GT/s Speeds

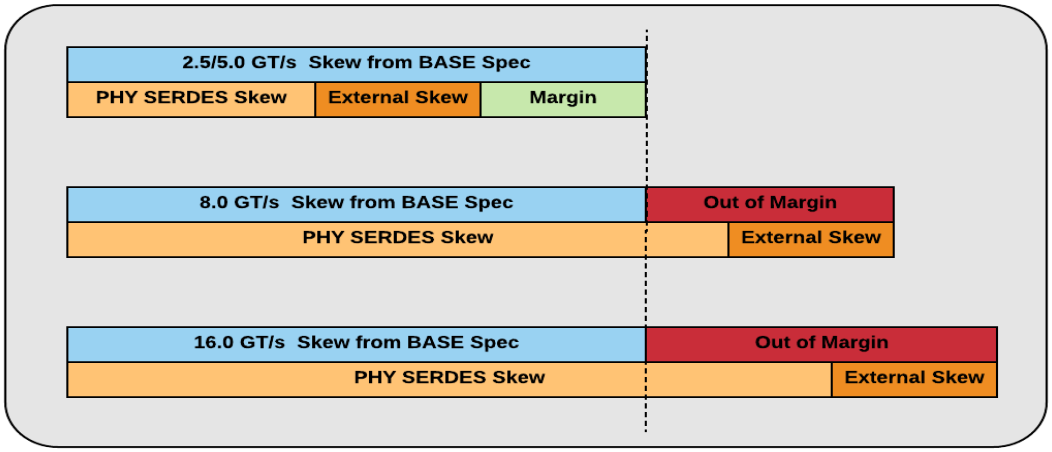

Lane-Lane De-skew in Numbers

- Skew Numbers defined by Base Spec at Serial Lines, and PHY SERDES, which gets accumulated at MAC I/O

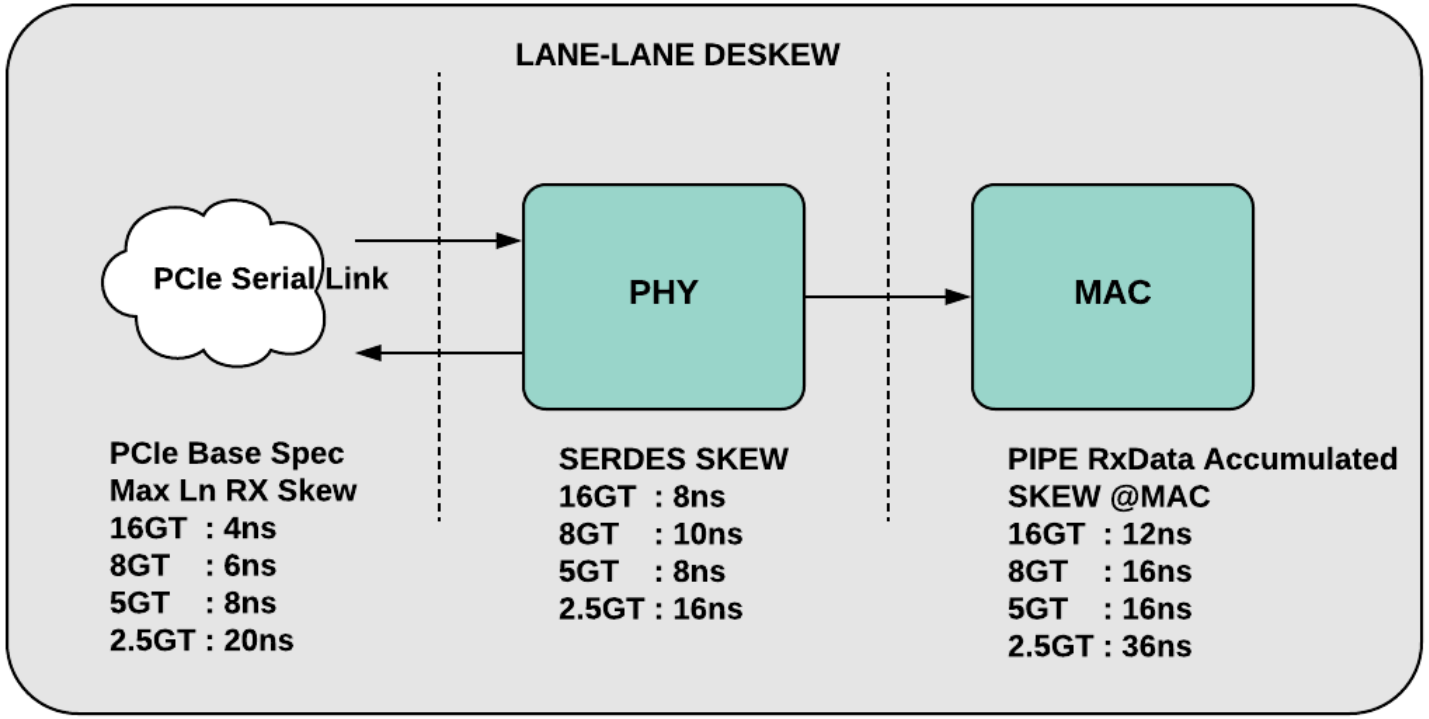

How to Resolve MAX Lane-Lane De-skew?

- Worst case Skew Numbers at 8.0GT/s & 16.0GT/s, needs an Extended De-skew FIFO.

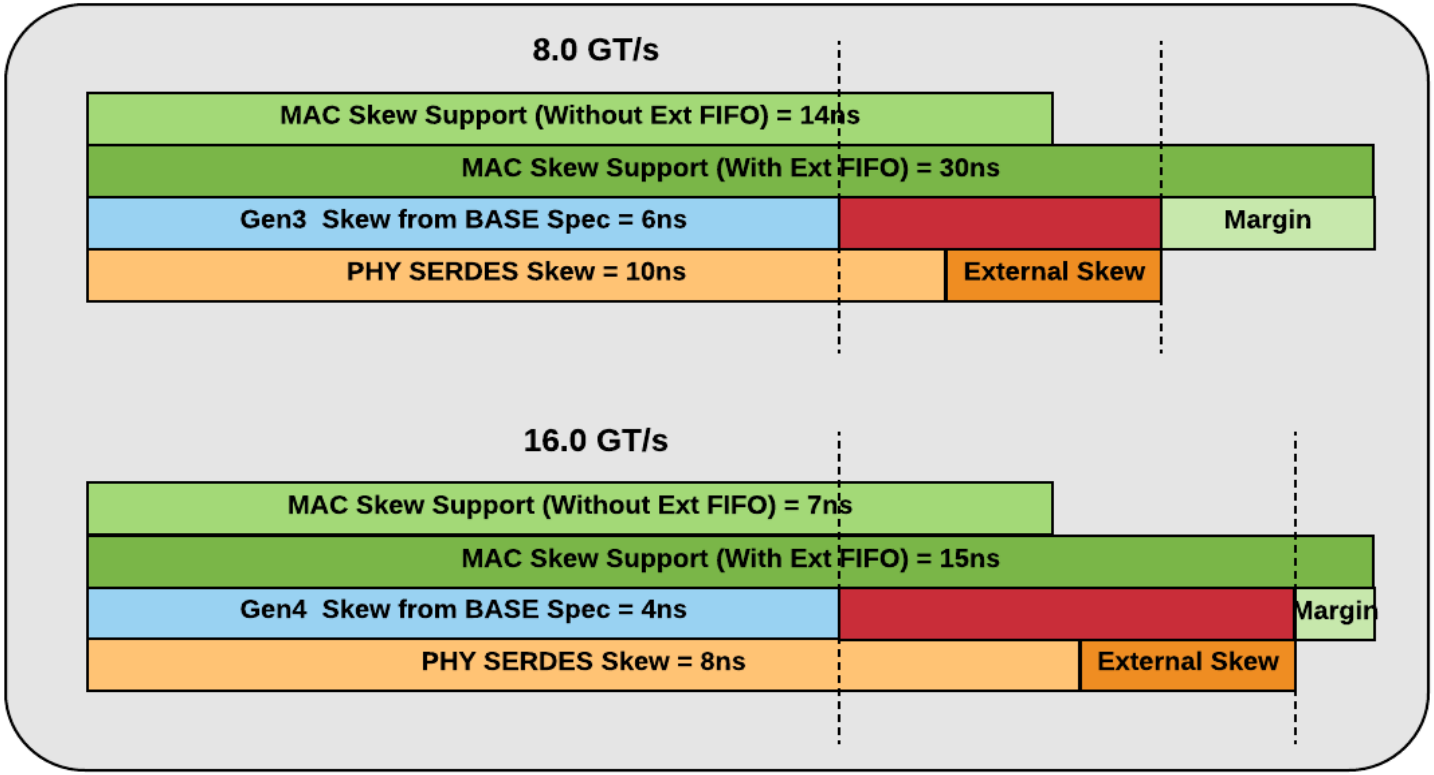

How to Resolve MAX Lane-Lane De-skew?

- Design an Extended DESKEW FIFO option to support excess Skew
- 8.0GT/s Skew requirement is 16ns and Extended Deskew FIFO provides MAC to manage 30ns skew
- 16.0GT/s Skew requirement is 12ns and Extended Deskew FIFO provides MAC to manage 15ns skew
- The disadvantage with Extended De-skew is it needs few more memory bits

	Without Extended FIFO	With Extended FIFO
MAX Skew 8.0GT	14 ns	30 ns
MAX Skew 16.0GT	7 ns	15 ns
Deskew FIFO Memory Bits	160	320

Rx Valid signal issue in L0s (P0s)

- Rx valid High and Electrical idle Low indicates the received data has valid data sequence

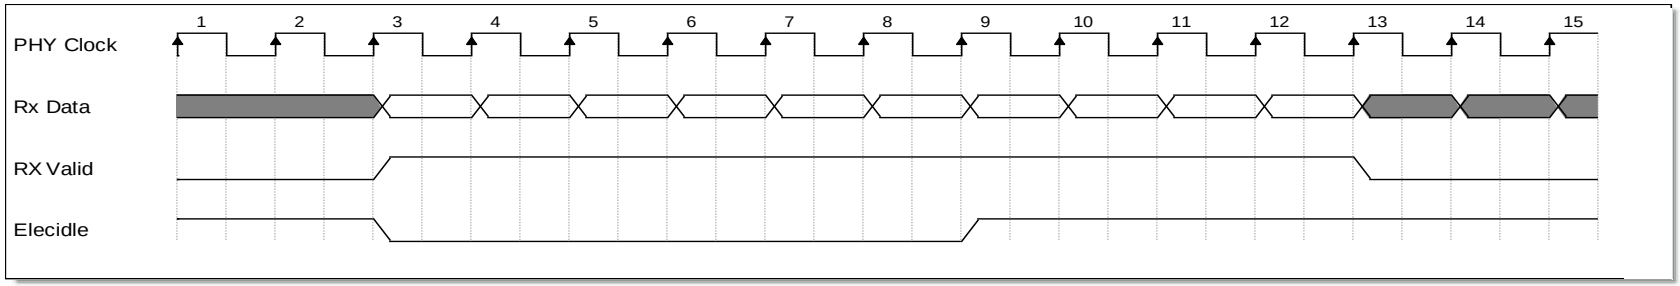

- We often see an issue while PHY is exiting from L0s, this Rx Valid from PHY is seldom going Low in the middle of data sequence.

- This causes alignment issue and can result in a different de-skew behavior across the lanes.

How to solve this scenario?

Solution to Rx Valid Issue in L0s (P0s)

- Rx Valid Filtering is designed to solve the issue. So other symbols are filtered out for the alignment
- MAC in addition to COM, should also detect other 8b/10b special symbols that are part of ordering sets (except TS). The result is that only SKP and TS are now used for the purpose of lane-lane alignment.

Rx Valid Filter

- BAD RxL0s Exit:

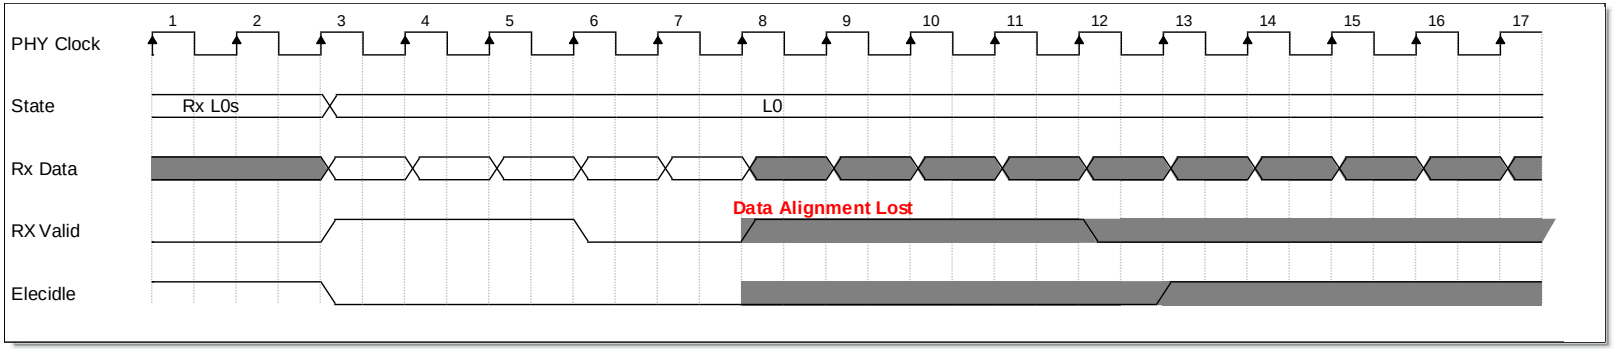

- GOOD RxL0s Exit:

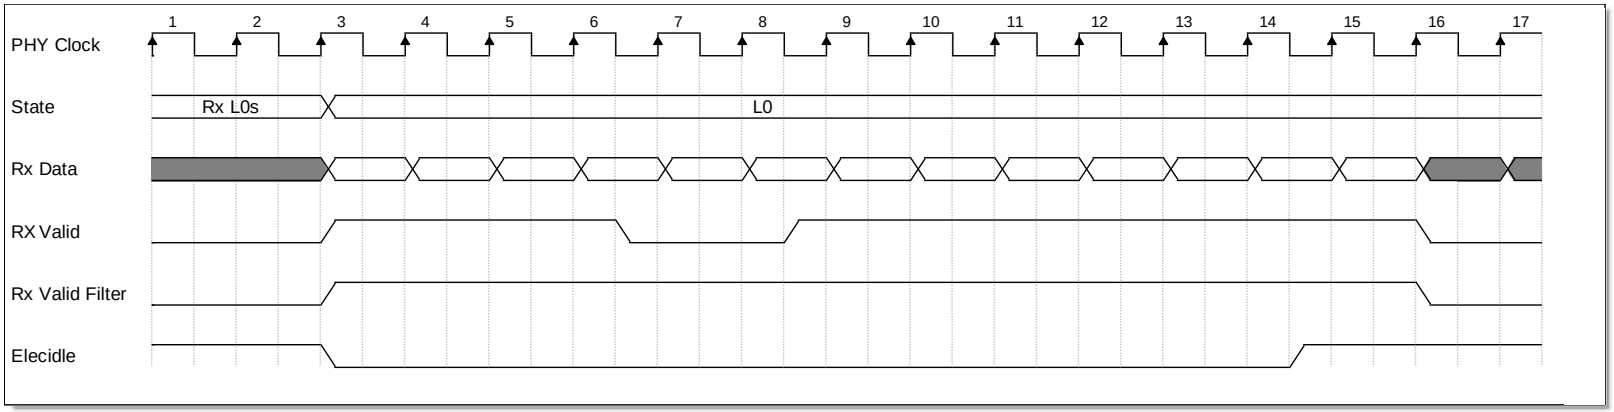

Rx Valid Filter

- Notes :
  - This features is enabled or disabled using a static switch
  - When the PHY doesn’t has Rx Valid issue this feature needs to be disabled

N_FTS (Number of Fast Training Sequences)

- PHY supporting Low Power modes, has a proprietary value for Number of Fast Training Sequences (N_FTS) to redo fast training while exiting low power state.

- Programming the wrong value may sometimes prevent the device to exit from Low power to L0 and instead causes link retraining.

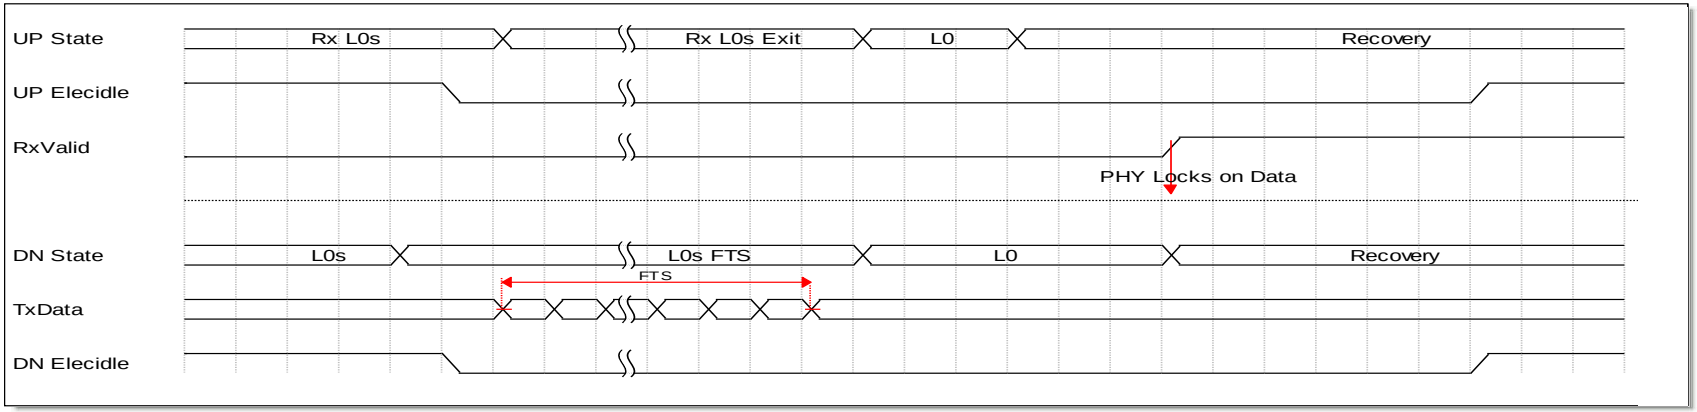

N_FTS (Number of Fast Training Sequences)

- Each PHY has specific N_FTS value separate for each Speed (2.5GT/s, 5.0GT/s, 8.0GT/s, 16.0GT/s).

- To ensure proper operation MAC should be Programmed with right values of N_FTS as defined in the PHY specification Document

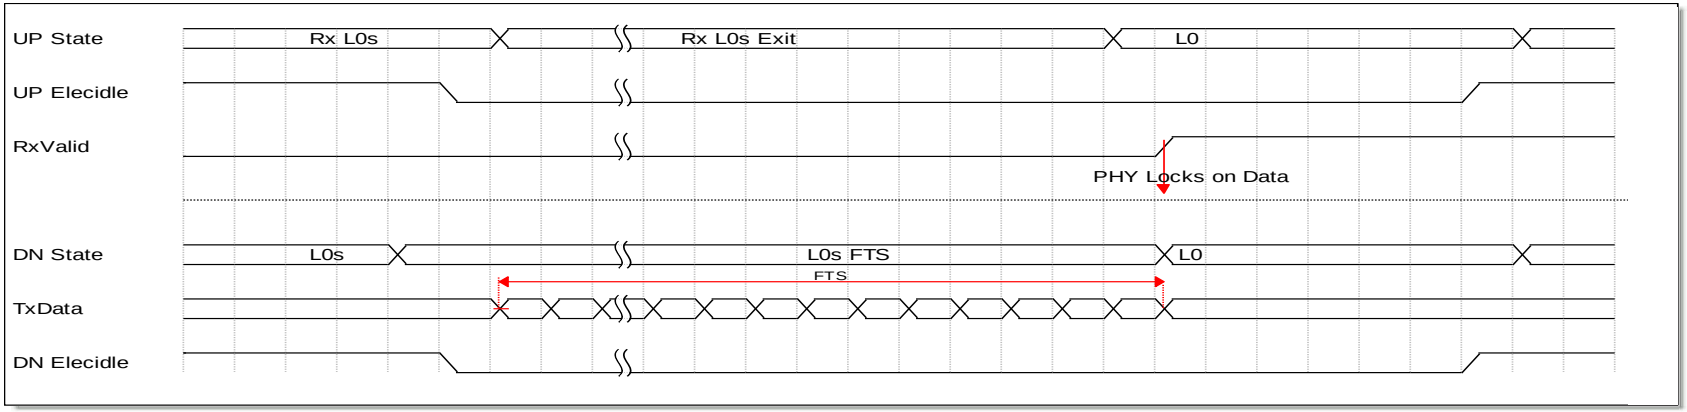

PHY Clock Generation

- PHY may generate the PHY Clock internally or use an External PLL for this purpose
- The CLK Rate is defined by the function of PCIe link rate and PHY interface width
- When PHY has the internal PLL and speed changes, internal logic will automatically generate the rated clock which is the output of PHY and input to MAC
- When PHY needs an external PLL and speed changes, Then How to handle the PHY Clock Change?

PHY Clock with External PLL

- PHY with External PLL needs handshake with MAC and PLL
- PHYClock_CHANGE_OK and PHYClock_CHANGE_ACK are the handshake signals for this purpose

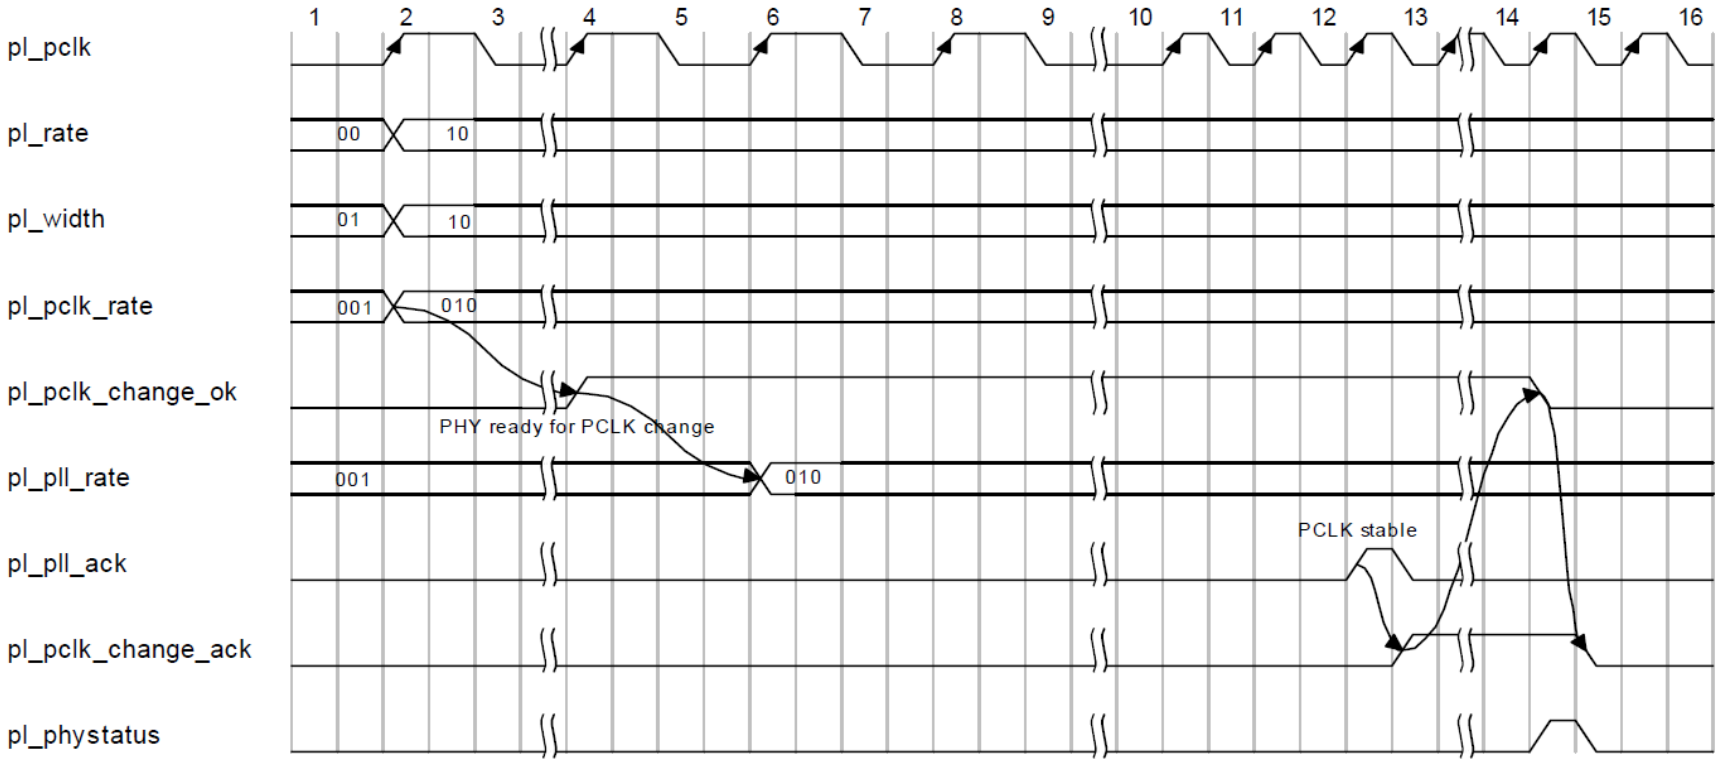

PHY Clock with External PLL

- When PLL is External to PHY, there is one more aspect to be taken care of is the PLL – PHY PCS and PLL – MAC clock path.
- Placement issues can cause clock skew and hence causing bit flips during enumeration and different alignment.
- PLL needs to be placed at the shortest distance from Mac and PHY PCS

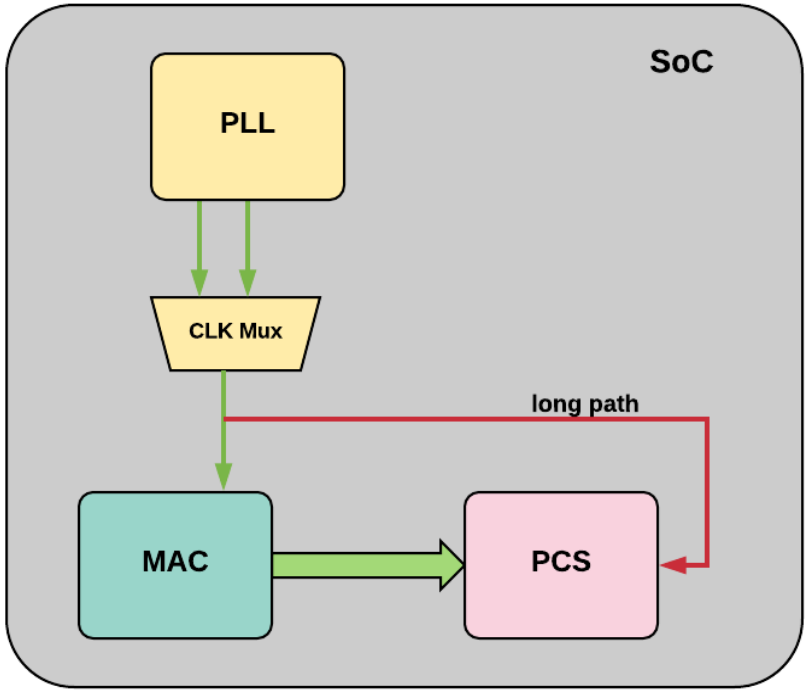

L1PM Substates

- Each PHY has its own mechanism using a signal for Power down or sideband signals handshake for L1PM Entry & Exit.
- What is the effect?
- How can we Handle this?

Value	Power down				Power State
PS0	0	0	0	0	L0/P0 State
PS1	0	0	0	1	L0s/P0s State
PS2	0	0	1	0	L1/P1 State
PS3	0	0	1	1	L2/P2 State
PS4	0	1	0	0	(PS4) PHY Specific
PS5	0	1	0	1	(PS5) PHY Specific
PS6	0	1	1	0	(PS6) PHY Specific
PS7	0	1	1	1	(PS7) PHY Specific
PS8	1	0	0	0	(PS8) PHY Specific
PS9	1	0	0	1	(P94) PHY Specific
PS10	1	0	1	0	(PS10) PHY Specific
PS11	1	0	1	1	(PS11) PHY Specific
PS12	1	1	0	0	(PS12) PHY Specific
PS13	1	1	0	1	(PS13) PHY Specific
PS14	1	1	1	0	(PS14) PHY Specific
PS15	1	1	1	1	(PS15) PHY Specific

L1PM Substates

- The Limitation !!

  A specific PHY use PS5 for L1.1 Entry and PS13 for L1.2 Entry,

  While another PHY use PS2 for L1.1/L1.2 Entry

  and other PHY use PS4 for L1.1 Entry and PS5 for L1.2 Entry

- Proposed Solutions !!
  - Designer Needs an external Glue Logic to Encode / Decode L1PM transitions on powerdown signal of MAC
  - MAC should adopt a mode specific to each PHY to make it more PHY agnostic

L1PM Sideband

- Sideband Signals for L1PM Entry and Exit are different between PHYs
- The common solution is to implement the L1 Substates Request, Ack and Status ports to indicate the L1PM Substates Entry & Exit

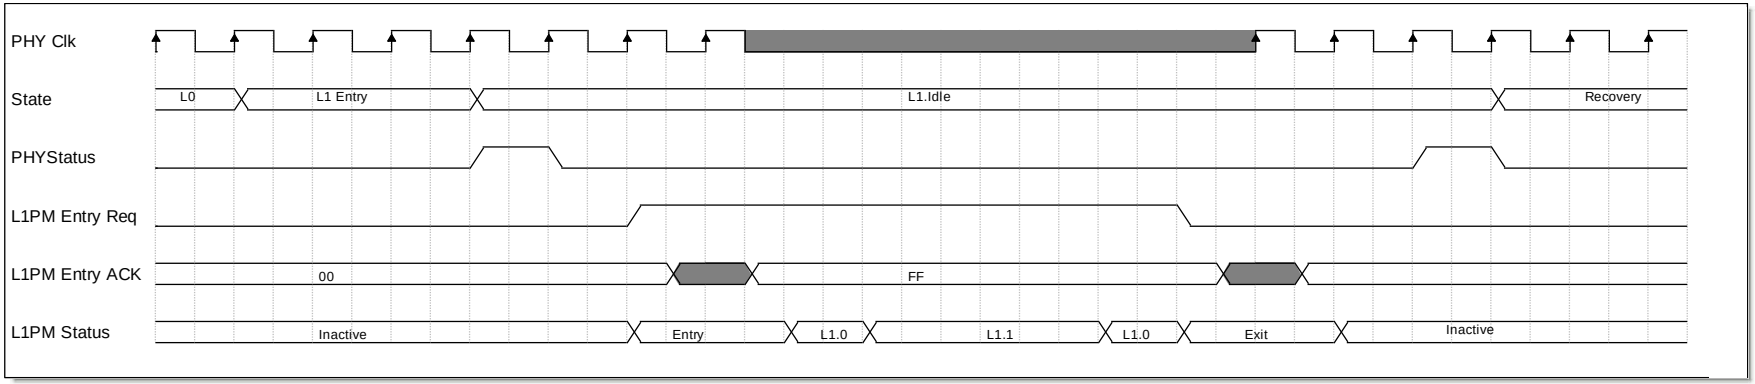

L1PM Sideband

Notes :

- The Sideband signals solution is suitable for most of the PHY’s
- User may need to tune these interface for some PHY’s

Lane Margining

- LMR Interface Issues between PHY's can be either due to
  - Different Interpretations by PHY and MAC designers

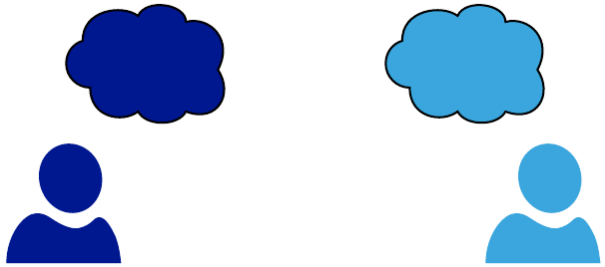

  - Supported Version of Specification

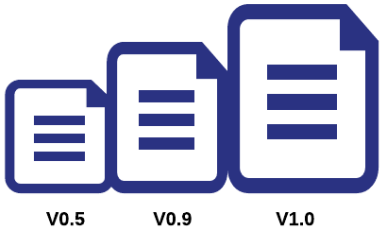

Lane Margining Issues Examples

- Different interpretation by PHY and MAC Designer
  - Sample Counter in PHY or MAC
  - PHY & MAC doesn’t support a transaction type
  - The parameters like Number of Timing steps, Maximum Timing Offset, Sampling Rate Timing are PHY Specific
- Supported Version and updates in the new Spec
  - PHY interface registers can be different by the version PHY and MAC
  - Speed change from 16.0GT/s to 8.0GT/s requires the MAC to wait for current Margining sequence to complete

Lane Margining Issues

- The objective of the Session is to cover the issues for PHY-MAC integration and so the Lane Margining topic is condensed.

For Elaborated LMR content, Please visit session

"**Challenges and Methods to Implement Lane Margining in a Heterogeneous System**"

Equalization Issues

- PHY will execute different steps of Equalization such as:
  - Coefficients Finetuning
  - Direction Change
  - Figure of Merit

Equalization - Coefficients Finetuning

- Coefficients Finetuning requires the list of coefficients exchanged between PHY and MAC.
- These Params (set of Coefficients) are queried by the PHY starting from the default.
- These Param Queries trip has a timeout (Query Timeout) and the best set of Params are used by PHY.
- What if PHY's doesn't respond till predefined Timeout? MAC needs a capability to increase this period or infinitely wait for completion followed by assertion of PHY status.

PHY Specific Options

- Power Management
  - BEACON, WAKE#
  - CLKREQ
- Multi-protocol PHY
  - Mode Select
- Simulation Mode
  - Receiver Detect
  - Equalization
  - Timeouts

Architecture Challenges - Multiple Controllers Sharing PHY

- Current Market trends demands Single Architecture Flexibility for multiple applications
- The Low power and cost sensitive application needs smaller link with better power budget
- Higher end applications needs bigger link with bigger power budget
- The un-used MAC and PHY instances are statically powered off
- The Sideband signals and Mainband signals needs to be carefully routed as per the Bifurcation configuration

Multiple controller sharing PHY (Port Bifurcation)

- PHY SERDES is per lane or per Quad. Lane is shared with MAC as a combination of x2, x4, x8 using the Bifurcation MUX.

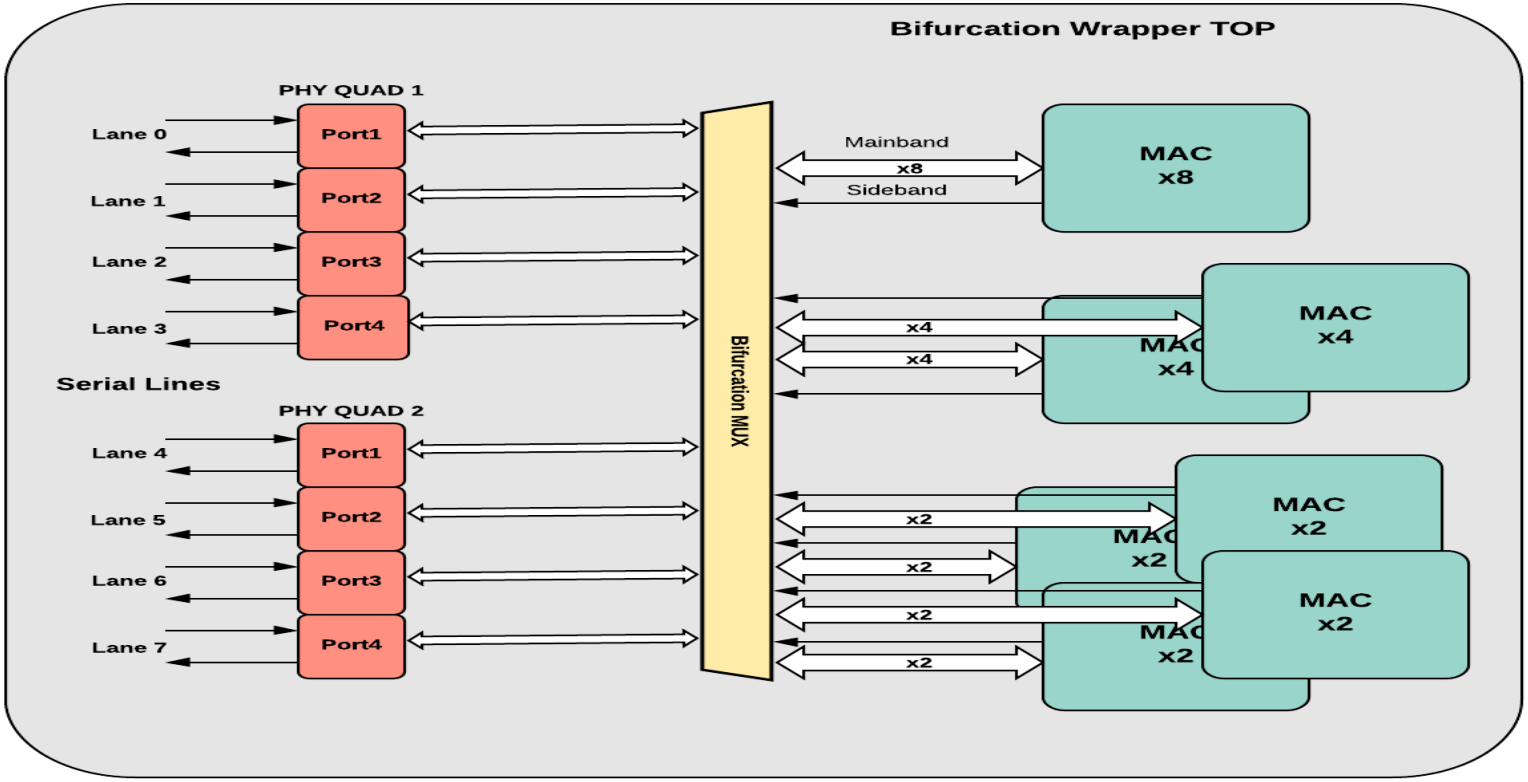

Multiple controller sharing PHY (Port Bifurcation)

- Advantages :
  - Architecture Flexibility – Single Architecture applies for multiple applications
  - Wide Market reach (More $$$ to Company)

- Drawbacks :
  - Multiple instances of Controllers adds complexity
  - Higher gate count
  - Extreme care needed for muxing Resets, Clocks, Sideband and Mainband
  - Verification Cycle

- Introduction
- Scenarios and Source of Issues
- Issues and Solutions During PCI Express PHY - Controller Integration
- **Conclusion**

Conclusion

- There are numerous pitfalls with PHY – MAC integration and having prior knowledge of these scenarios for potential issues will minimize the risk
- Solving listed issues b/w MAC – PHY enables Designers to conquer the Integration task
- The scenarios and recommendations in this presentation are based on 17+ years of experience integrating PCIe controllers with PHYs at speeds from 16.0GT/s to 2.5GT/s on various process nodes from 40nm to 7nm.

Thank you for attending the PCI-SIG Developers Conference 2018.

For more information, please go to www.pcisig.com

Don't forget to submit your feedback via the mobile app!

Download the **Crowd Compass** app and then search for **PCI-SIG Developers Conference** or entering the following URL into your mobile browser: https://crowd.cc/s/1rKy0
Enter event code: **DevCon2018**

Alternatively, access here: https://crowd.cc/pcisig2018

Note: Create an account within the app so Admin knows who to contact if selected as the prize winner.

Each session feedback is provided is equivalent to 1 raffle entry (up to 11 sessions).
General survey feedback = 1 raffle entry.